Developing EFSM-based stateful applications with FlowBlaze.p4 and ONOS

Daniele Moro, Davide Sanvito^, Antonio Capone**

* Politecnico di Milano, Italy
^ NEC Laboratories Europe, Germany

3rd P4 Workshop in Europe (EuroP4)
December 1, 2020
Barcelona (Spain) - Online Workshop

Introduction

5G and Mobile Edge Computing requires offloading of network functions to data plane

- **P4**: reference language for data plane programming
- **State Machines**: powerful abstraction to develop stateful packet processing
- **FlowBlaze [NSDI '19]**: EFSM-based stateful packet processing architecture

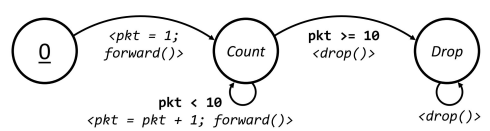

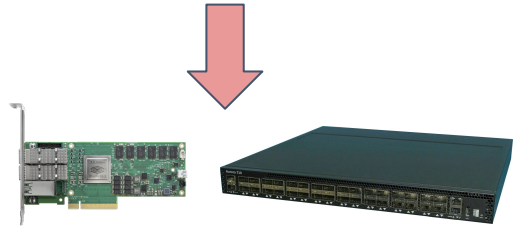

FlowBlaze currently lacks of:

- Prototyping platform and P4 implementation*
- Integration with DC-style fabric

*FlowBlaze.p4 [NFV-SDN '20]

FlowBlaze.p4*

- FlowBlaze library implementation in P4
- Open source library
- GUI to automatically translate EFSM into table entries
- Exploit all the tools from the P4 Community
- Targets: BMv2 and V1Model

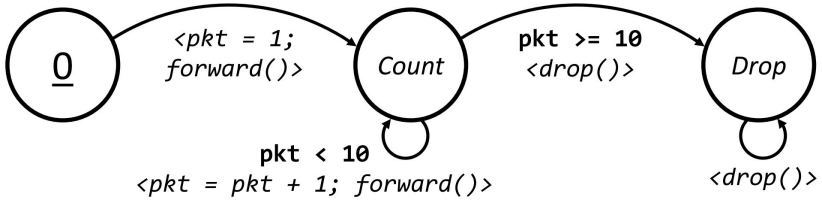

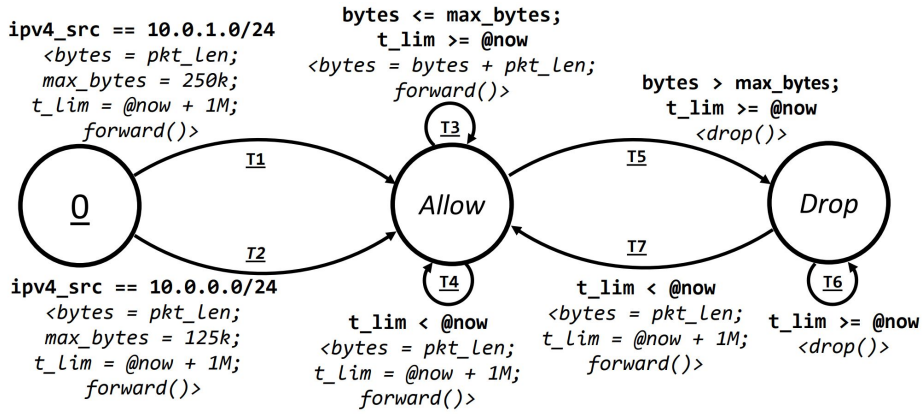

** D. Moro, et al. "**FlowBlaze.p4: a library for quick prototyping of stateful SDN applications in P4**" IEEE NFV-SDN 2020*

FlowBlaze.p4 - configuration

Compile-time configuration

VS

Run-time configuration

- Draw the EFSM
- Auto generate the configuration via Python backend

```
#include "../flowblaze_lib/flowblaze_metadata.p4"
#include "headers.p4"
#include "metadata.p4"
#include "../flowblaze_lib/flowblaze.p4"
 ...
#define FLOW_SCOPE { hdr.ipv4.srcAddr }
#define CUSTOM_ACTIONS_DEFINITION @name(".FlowBlaze.forward") \
                                  action forward() { \
                                    \
                                  } \
                                  @name(".FlowBlaze.drop") \
                                  action drop() { \
                                    mark_to_drop(standard_metadata); \
                                    exit; \
                                  }
#define CUSTOM_ACTIONS_DECLARATION forward; drop;
// Configuration parameter left black because not needed
//    #define METADATA_OPERATION_COND
//    #define EFSM_MATCH_FIELDS
//    #define CONTEXT_TABLE_SIZE
 ...
 apply {
     if (hdr.ethernet.isValid()) {
         FlowBlaze.apply(hdr, meta, standard_metadata);
         t_l2_fwd.apply();
     }
 }
```

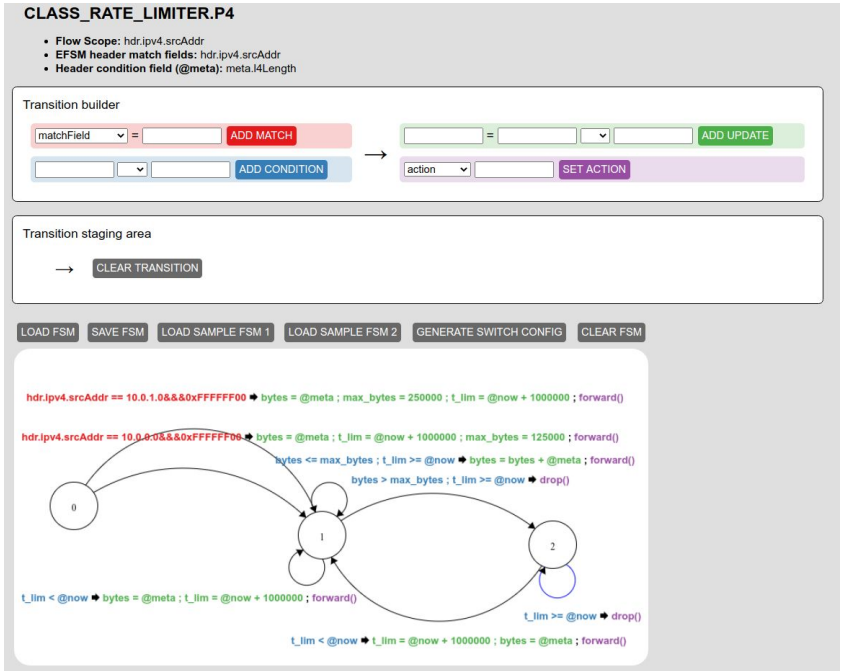

ANTLAB

ONOS integration

- Add `flowblaze.p4` library to `fabric.p4` pipeline
- Exploit current Trellis apps to program the rest of the fabric.p4 pipeline
  (routing, bridging, link discovery…)
- FlowBlaze ONOS app to control the FlowBlaze portion of the pipeline

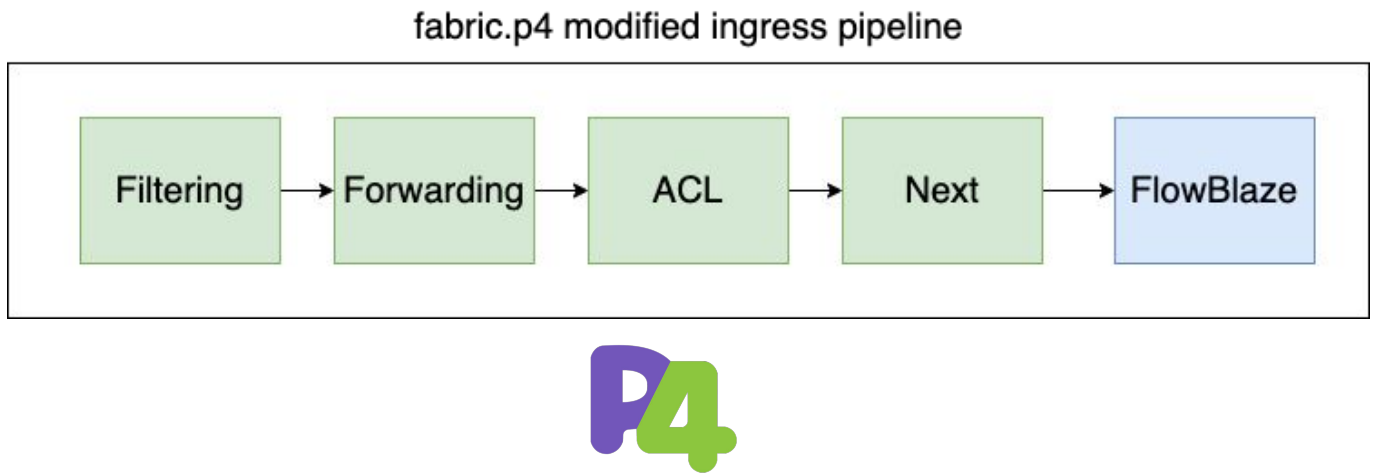

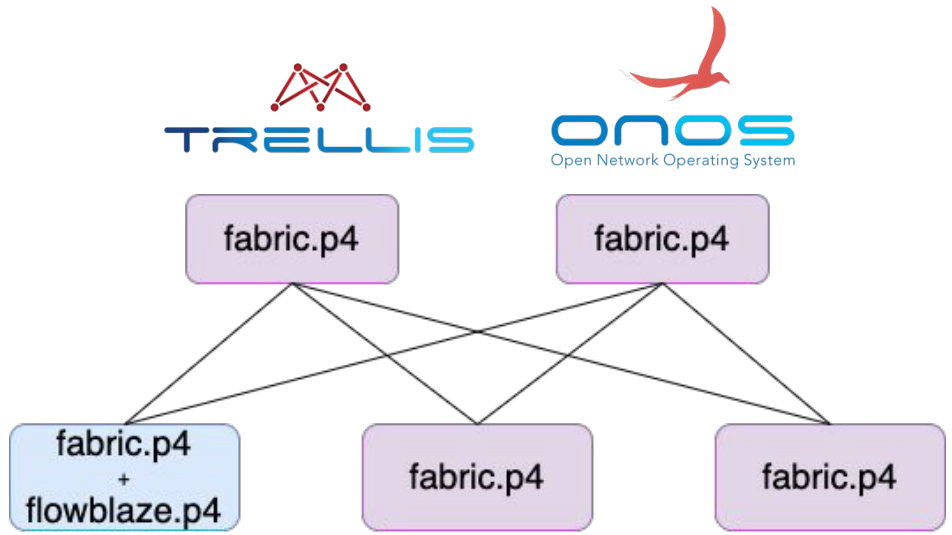

FlowBlaze.p4 + ONOS WorkFlow

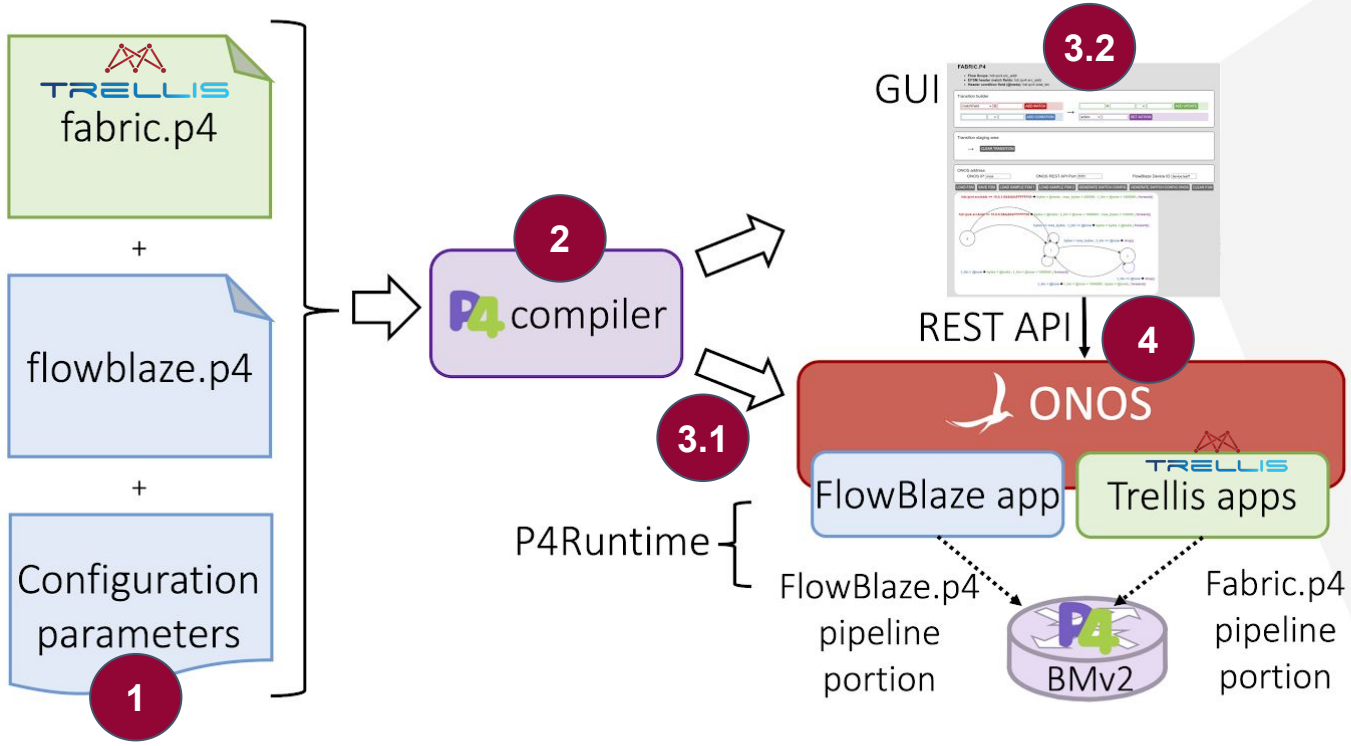

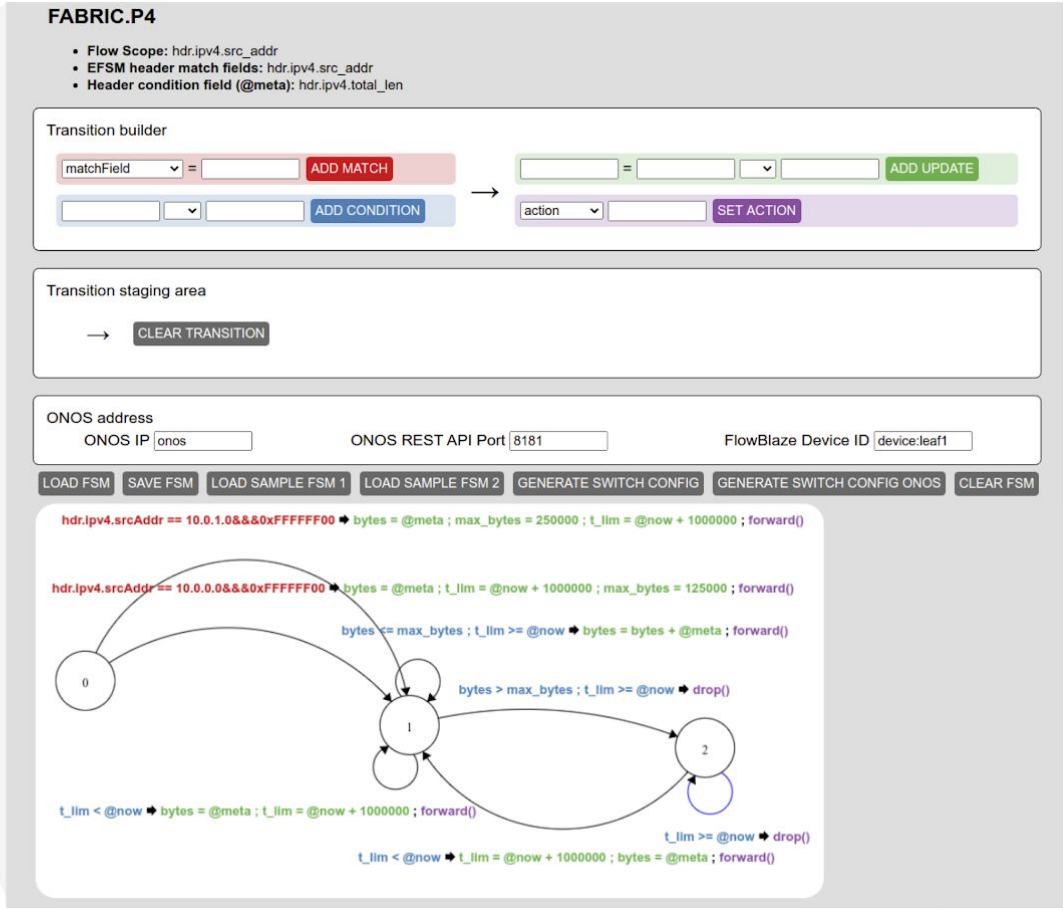

ANTLAB

Daniele Moro
daniele.moro@polimi.it

ANTLAB | ADVANCED NETWORK TECHNOLOGIES LAB

https://github.com/ANTLab-polimi/ONOS-flowblaze

ADVANCED
NETWORK
TECHNOLOGIES
LAB